

Finite Element Interpolation for Combined Classical/Quantum Mechanical Molecular Dynamics Simulations

CHRISTIAN D. BERWEGER, WILFRED F. VAN GUNSTEREN,
FLORIAN MÜLLER-PLATHE*

*Laboratorium für Physikalische Chemie, Eidgenössische Technische Hochschule Zürich, ETH Zentrum,
CH-8092 Zürich, Switzerland*

Received 7 October 1996; accepted 25 February 1997

ABSTRACT: A method is presented to interpolate the potential energy function for a part of a system consisting of a few degrees of freedom, such as a molecule in solution. The method is based on a modified finite element (FE) interpolation scheme. The aim is to save computer time when expensive methods such as quantum-chemical calculations are used to determine the potential energy function. The expensive calculations are only carried out if the molecule explores new unknown regions of the conformation space. If the molecule resides in regions previously explored, a cheap interpolation is performed instead of an expensive calculation, using known neighboring points. We report the interpolation techniques for the energies and the forces of the molecule, the handling of the FE mesh, and an application to a simple test example in molecular dynamics (MD) simulations. Good performance of the method was obtained (especially for MD simulations with a preceding Monte Carlo mesh generation) without losing accuracy. © 1997 John Wiley & Sons, Inc.
J Comput Chem **18**: 1484–1495, 1997

Keywords: energy hypersurface interpolation; force interpolation; simplex; molecular dynamics algorithm; Monte Carlo algorithm

* *Present address:* Max-Planck-Institut für Polymerforschung,
Ackermannweg 10, D-55128 Mainz, Germany.

Correspondence to: W. F. van Gunsteren.

Introduction

The applicability of combined quantum mechanical/molecular mechanical (QM/MM) models (hybrid models)^{1,2} is limited by the computational cost of the QM part. *Ab initio* calculations are too expensive for larger systems, while the cheaper semiempirical models are often too inaccurate and unreliable to be useful.^{3,4}

The aim of this article was to develop an interpolation scheme that replaces the expensive quantum chemical calculations of points of the potential energy surface by a much cheaper interpolation wherever possible. Points of the potential energy surface are calculated by quantum mechanics if the region of the conformation space is visited for the first time. When the molecule comes back to a region of the conformation space that is already known, an interpolation (cheap) between known points nearby is performed instead of an explicit calculation (expensive) at the point of interest.

The method is evaluated using a simple model system of an isomerizing molecule in solution. Its computational efficiency is tested using a classical force field rather than quantum chemistry at the current stage of development.

During the simulation, points of the potential energy surface are calculated (energy plus energy gradient). Because of the expense of these calculations, this information is not thrown away after a molecular dynamics (MD) step but is kept in the memory instead. After some time, enough mesh points (vertices) will be available to form simple geometric objects (the finite elements, FEs) consisting of $n + 1$ vertices, n being the number of dimensions of the conformation space. These convex objects are taken to be triangles in 2 dimensions, tetrahedra in 3 dimensions, or their n -dimensional generalization, the so-called n simplices in n dimensions.

When a new point in the conformation space is reached in a simulation step, whether this point lies within an existing element is evaluated. If it does, the potential energy and its gradient, the negative forces, are obtained from an interpolation. If not, they are calculated. The new point is then added as a new vertex to the mesh of known vertices and attempts are made to establish new elements.

The method is expected to be useful and applicable if the following conditions are satisfied.

- The calculations of the potential energy surface are computationally much more expensive than the interpolations. If sophisticated quantum-mechanical methods are used, this will be the case.
- Only a few degrees of freedom describe the interesting (and expensive) part of the system (low dimensionality). The others can either be frozen by constraints or are treated by a cheap classical force field^{5,6} (e.g., the solvent).
- The interesting degrees of freedom are bounded and nondiffusive (e.g., internal coordinates of the solute molecule). If the space of only a few dihedral angles and bond angles is considered, it is rather small and even circular.
- There are small regions of conformation space where the molecule will often reside (nondiffusive system; few, sharp energy minima; few reaction pathways).
- The simulation time is long enough for the system to revisit previously explored regions of the conformation space. The longer a simulation lasts, the more efficient the method is expected to be.

With a view to practical applications, the most critical point is the restriction to a few relevant dimensions. The danger is that possible pathways are not included among the few degrees of freedom that are allowed in the simulation. In general, it cannot be assumed *a priori* that a system will react along only one specific pathway.

The method was successfully applied in the 2-dimensional case.⁷ In this work it was extended to multiple dimensions but was so far only tested and applied in 3 dimensions.

Method

Our method makes use of ideas from FE interpolation methods. However, many of the basic aspects of FE must be modified. For a comprehensive FE theory, see, for example, refs. 8 and 9.

TRIANGULATION AND ELEMENT SHAPES

Because simplicial FEs seem to be the easiest to handle, this kind of FE was chosen rather than rectangular FEs. In simplices (n -dimensional tetrahedra), the corresponding interpolation polynomial

als are easy to derive. They are also considered to lead to earlier interpolations in an MD simulation because fewer vertices are needed to form an element.

The main difference from FEs methods as used traditionally is the triangulation, which is not known in advance but is progressively built up, point by point, during the simulation. Thus, it is impossible to select an optimum grid with elements of regular shape. On the contrary, the mesh points have to be taken as they come from the simulation. Hence, elements that violate the triangulation conditions or otherwise have poor properties have to be intercepted and rejected.

In a correctly triangulated domain, and using a linear or quadratic polynomial for interpolation, the interpolant (e.g., the energy) is continuous. If the triangulation is incorrect (e.g., due to intersecting elements that cover part of the space twice) or if the faces of two neighboring elements do not fully match, the continuity in energy is not achieved. In addition, there may be ambiguity if a part of the conformation space is covered by multiple elements, depending on which element is used for the interpolation. Furthermore, it is advantageous to omit elements with very acute or very obtuse angles, because they often lead to inaccurate interpolation. Moreover, the interpolant becomes badly behaved in regions where many elements share the same vertex.

INTERPOLATION OF ENERGY AND FORCES

We employ a second-order polynomial for the interpolation of the energy. This leads to an interpolant that is continuous in energy, but the gradients will only approximately match over the elements' boundaries. The example is given for 3 dimensions but is easily generalized. The conformation space is spanned by three coordinates r_1 , r_2 , and r_3 (dihedral angles in our test system). The energy interpolation function is

$$E = C_1 + r_1 C_2 + r_1^2 C_3 + r_2 C_4 + r_2^2 C_5 + r_1 r_2 C_6 + r_3 C_7 + r_3^2 C_8 + r_1 r_3 C_9 + r_2 r_3 C_{10}. \quad (1)$$

The polynomial has 10 unknown coefficients C_1 – C_{10} that have to be determined from information at the vertices. Using the energies and gradients, which are needed for the MD part of the simulation, yields 16 knowns (one energy and three gradient components at four vertices). The information is reduced in the following way. Using the energies and gradients at every pair of vertices i

and j at the positions \mathbf{p}_i and \mathbf{p}_j , the energy E_m of the midpoint of the connecting edge can be estimated by a third-order polynomial, using the energies E_i and E_j and directional derivatives $(\partial E / \partial p)_{\mathbf{p}_i}$ and $(\partial E / \partial p)_{\mathbf{p}_j}$ along the edge $\mathbf{p} = \mathbf{p}_i - \mathbf{p}_j$.

$$E_m = \frac{1}{8} \left(\left(\frac{\partial E}{\partial p} \right)_{\mathbf{p}_j} - \left(\frac{\partial E}{\partial p} \right)_{\mathbf{p}_i} \right) |\mathbf{p}_i - \mathbf{p}_j| + \frac{1}{2} (E_i + E_j). \quad (2)$$

Using the energies of the vertices and all of their edge midpoints, there is now just enough information to determine a quadratic polynomial for interpolation: four vertices plus six midpoints determine the 10 unknowns in eq. (1). This technique can easily be generalized to more dimensions.

MASTER ELEMENTS

It is common in FE methods to use shape functions with a special property. A shape function takes the value of unity at one of the mesh points (a vertex or an edge midpoint) and is zero at all others. The local interpolant is then constructed as a linear combination of shape functions, using the energies of the mesh points as coefficients.

If the elements are irregularly spread over the space, they all have distinct shape functions. However, it is possible to transform the irregular (triclinc) element into a master element (Cartesian). The master element is a special virtual element, which has the origin and unity on every axis as vertices. The shape functions of the master element are calculated once, prior to any interpolation. For the 3-dimensional case, the 10 quadratic shape functions Ψ for FE interpolation read

$$\begin{aligned} \Psi_{(0,0,0)}(r_1, r_2, r_3) &= 1 - 3r_1 + 2r_1^2 - 3r_2 + 2r_2^2 - 3r_3 \\ &\quad + 2r_3^2 + 4r_1 r_2 + 4r_1 r_3 + 4r_2 r_3, \end{aligned}$$

$$\Psi_{(1,0,0)}(r_1, r_2, r_3) = -r_1 + 2r_1^2,$$

$$\Psi_{(0,1,0)}(r_1, r_2, r_3) = -r_2 + 2r_2^2,$$

$$\Psi_{(0,0,1)}(r_1, r_2, r_3) = -r_3 + 2r_3^2,$$

$$\Psi_{(1/2,0,0)}(r_1, r_2, r_3) = 4r_1 - 4r_1^2 - 4r_1 r_2 - 4r_1 r_3,$$

$$\Psi_{(0,1/2,0)}(r_1, r_2, r_3) = 4r_2 - 4r_2^2 - 4r_1 r_2 - 4r_2 r_3,$$

$$\begin{aligned}
\Psi_{(0,0,1/2)}(r_1, r_2, r_3) &= 4r_3 - 4r_3^2 - 4r_1r_3 - 4r_2r_3, \\
\Psi_{(1/2,1/2,0)}(r_1, r_2, r_3) &= 4r_1r_2, \\
\Psi_{(1/2,0,1/2)}(r_1, r_2, r_3) &= 4r_1r_3, \\
\Psi_{(0,1/2,1/2)}(r_1, r_2, r_3) &= 4r_2r_3,
\end{aligned} \quad (3)$$

where the index of Ψ denotes at which point of the master element the shape function has the value of unity.

Any irregular FE can be transformed into the master element by a coordinate transformation (translation, followed by a shear rotation). The transformation of the point \mathbf{p}_t subjected to interpolation is achieved by the matrix operation

$$\mathbf{p}_M = \mathbf{M}^{-1}(\mathbf{p}_t - \mathbf{p}_0), \quad (4)$$

with the transformation matrix \mathbf{M} defined by the translated vertices of the triclinic element,

$$\mathbf{M} = ((\mathbf{p}_1 - \mathbf{p}_0)(\mathbf{p}_2 - \mathbf{p}_0)(\mathbf{p}_3 - \mathbf{p}_0)), \quad (5)$$

where the \mathbf{p}_0 to \mathbf{p}_3 are the positions of the vertices of the triclinic element. The point \mathbf{p}_M then has the same relative position in the master element as the point \mathbf{p}_t has in the triclinic element. The energy and the gradient of the transformed point are then calculated using the interpolation polynomial. The energy can be taken directly from the master element because it is invariant under coordinate transformation. However, the gradient requires backtransformation to the irregular element. This is achieved by the matrix operation

$$\mathbf{g}_t = (\mathbf{M}^T)^{-1} \mathbf{g}_M. \quad (6)$$

BUILDING UP THE MESH

Finding the Element that Contains the New Point

To find the element that contains the new point, the known vertices are scanned. Only vertices that are nearer than a certain cutoff distance (maximum element edge length) are considered and added to a list of neighbors. As soon as a vertex is part of any element, it is checked whether the point lies within it or not. If it does, the desired element is found and an interpolation is carried out. If not, the element that lies on the opposite side of the face of the current element, of which the point is found to lie outside, if it exists, is checked next. Moreover, all elements already examined are flagged to prevent a second evaluation. If no appropriate element is found, a complete list of neighboring vertices is established

instead. The energy and its derivatives are then calculated rather than interpolated.

Adding Elements to the Triangulation

Using the list of neighboring vertices, all possible combinations of vertices that may form new elements with the new point are generated and examined whether a valid and well-shaped element is formed. However, the number of possible combinations rises dramatically with the number of neighbors. Because the element validity checks are not straightforward, the evaluation of the large amount of possible elements may become slow and most of the proposals will be rejected anyway. So it would be advantageous to reduce the number of neighbors by removing those that will not form well-shaped elements. Hence, the following filters are applied to reduce the number of neighboring vertices, the number of proposed elements, the number of new elements that are ultimately added to the triangulation, and to optimize the overall triangulation.

Remapping. The space spanned by dihedral angles is actually cyclic (periodic). This property was not implemented in the current version. Instead, undesirable consequences are circumvented with little effort: the borders of the conformation space actually considered are shifted in such a way that they come to lie in regions of high energy where few points are encountered and probably no elements can be installed. Otherwise, if the borders lie in densely populated regions, there would be many elements prevented from extending over the borders. This technique was found to be very effective in practice. However, it is only applicable if the regions of high energy are known in advance. This is the case in the current test system.

Shielding. In MD simulations the trajectory often forms a more or less straight line that leads to elements that will either have bad angles or form narrow tubes along the trajectory path, preventing better elements from being built later. Shielding means that of the neighbors that are less than a certain angle apart (15°), viewed from the central (new) point, only the nearest can be seen. Thus, a neighbor is interpreted as a diffuse cloud that is able to shield other neighbors behind it. The shielded points are then removed from the neighbor list. In general, the reduction of neighbors makes the combinatorics to form elements more efficient. The number of neighbors is heavily

reduced by this relatively simple criterion. As a consequence, the number of combinations that may span elements becomes dramatically smaller.

Shadowing. Shadowing consists of removing all the neighbors that lie behind an existing or newly constructed element, viewed from the new point. Such points cannot be connected to form an edge. The concept of shadowing is needed to avoid triangulation violations such as intersecting elements.

Distance Ordering. An ordering of the neighbors with increasing distance from the new point is needed to prevent big elements from being built that enclose nearer points. After sorting the neighbors, it is guaranteed that smaller elements are established, if successfully passing through the other restrictions, before larger ones are considered.

Bad Angle Rejection. Elements with very acute ($< 15^\circ$) or very obtuse angles ($> 165^\circ$) are neither suited for interpolation nor for the handling of the elements. Thus, such bad-shaped elements have to be intercepted. Unfortunately, a simplex has many angles that all have to be tested, except the ones involving the new vertex as the middle point, which were already eliminated by the shielding procedure.

Triangulation Consistency Checks. In rare cases it is possible that some neighbors that would lead to a wrong triangulation are not eliminated by the foregoing filters. Thus, another triangulation consistency check is needed. For that reason, the fact is used that a face of a simplex, defined by n of its vertices, always has two distinct sides that can be occupied by two neighboring elements. Whether the element being built up would lie on the same side as an existing element is examined. If so the element is rejected.

Gradient Accuracy Checks. Unlike the energy, the gradient is not continuous over the elements' boundaries using a quadratic polynomial for interpolation. To make sure that the gradient discontinuity is not too large, the new element is rejected if, at any vertex, the deviation of the interpolated gradient from the one explicitly calculated (when the vertex was added) exceeds a certain limit, here 1% of the maximum force. This filter allows the use of larger elements without making the interpolations less accurate.

There may be redundancy in applying all these filters. However, it has been found that every filter is necessary to obtain a correct triangulation and

is useful to obtain an efficient and accurate interpolation.

Results and Discussion

MODEL SYSTEM "HEXA-ARGANE"

A very simple molecular system was chosen to test the method. It consists of a chain of argon atoms in a solution of fluid argon. The "argane" chain has three dihedral angles, indicated by r_1 , r_2 , and r_3 , that span the conformation space for the FE interpolation. All other internal degrees of freedom, such as bond lengths and bond angles, are held fixed by constraints. The dynamics of the molecule is governed by arbitrarily chosen potential energy functions imposed on the dihedral angles.

All atoms (solute and solvent) are argon atoms with mass $m = 39.948$ amu and Lennard-Jones parameters $\varepsilon = 0.996$ kJ/mol and $\sigma = 0.341$ nm. All bond lengths are constrained¹⁰ to 0.148 nm, bond angles to 120° . The potential energy function imposed on the dihedral angles r_i reads $E(r_i) = 2.5$ kJ/mol $\cdot (1 - \cos(2r_i))$ on every dihedral, which means a two-well potential energy function (minima 0° , cis, and 180° , trans) with a barrier of 5 kJ/mol in between. No other intramolecular interactions are taken into account.

The following parameters were used in all simulations unless explicitly indicated. A periodic cubic box of 2.15-nm edge length was used, containing 216 solvent argon atoms and one hexa-argane molecule (density 1.4818 g/cm³). The leapfrog integration time step was 0.01 ps. Non-bonded forces were cut off at 1 nm distance. For NVT simulations, the system was weakly coupled to a temperature bath at 300 K with 1 ps relaxation time.¹¹ Two hundred thousand MD steps were performed, equivalent to a total simulation time of 2 ns. The relative geometric tolerance in the SHAKE¹⁰ procedure was set to 10^{-6} .

SIMULATIONS IN NVE ENSEMBLE

Simulations in the NVE ensemble were performed to examine the energy conservation using the interpolation scheme. Moreover, it was checked whether the interpolation would significantly change the dynamic and static properties of the isomerization. A reference simulation without interpolations was performed.

The simulations with interpolations all lost more energy than the reference simulation did (Table I).

However, no trend was visible among the simulations using different element sizes. This is in strong contrast to simulations without applying the gradient accuracy check (not shown), in which the drift increased dramatically with the larger elements.

To get the correlation time (relaxation time) of the dihedrals remaining in either the cis or the trans conformation, each of the three dihedral angles was mapped to a state function that is 1 for a cis configuration ($|r_i| < 90^\circ$), or 0 for a trans angle ($|r_i| \geq 90^\circ$). The autocorrelation of this function can be well fitted by a stretched exponential function, $\exp(-(t/\alpha)^\beta)$. The integral of this function is analytically available and determines the relaxation time $\tau = \int_0^\infty \exp(-(t/\alpha)^\beta) dt = \alpha/\beta \Gamma(1/\beta)$.

To obtain the uncertainty of the cis fraction and the correlation time, the trajectory was split into five parts and each of these was analyzed by the same procedure as the entire trajectory (jackknife method¹²). The standard deviation of these five values is indicated in Table I and serves as a measure of reliability of the corresponding average. Figure 1 shows that the correlation times for various element sizes differ from each other within

the error bars. The two outer dihedrals should yield equal values due to symmetry. This is approximately achieved.

PERFORMANCE AND ACCURACY OF FE INTERPOLATION

Different maximum element sizes were used to investigate the performance (number of interpolations per MD step, ratio of interpolations to explicit calculations) and the accuracy (deviation of the interpolated to the calculated value) of the interpolation. To check the accuracy all interpolations were compared to the true value which was calculated in addition to the interpolation. If the deviation between calculated and interpolated gradient exceeded the limit of 0.05 kJ/(mol rad), this value was identified as inaccurate. Inaccurate gradients occurred in larger numbers with larger elements. The employed quadratic interpolation polynomial did not fit the cosine type potential energy surface of the dihedral angle well if the mesh points were far apart. Yet, the interpolated energies were mostly very accurate. The maximum

TABLE I. Energetics and Static and Dynamic Properties of Hexa-Argane Test System, Using Different Maximum Edge Lengths of Elements.

Max. Simplex Edge Length (rad)	Total Energy (kJ / mol)			cis Fraction		Correlation Time τ (ps)	
	Average	RMS Fluctuation	Drift	Average	SD	Average	SD
0.0	-122.5	0.31	-0.36	0.60	0.063	34	11
				0.55	0.048	63	31
				0.60	0.091	35	8
0.09	-122.8	0.48	-1.35	0.62	0.040	37	12
				0.65	0.061	55	24
				0.65	0.048	33	6
0.15	-122.6	0.42	-1.06	0.67	0.034	37	18
				0.68	0.065	52	16
				0.66	0.050	41	22
0.30	-122.8	0.46	-1.26	0.55	0.066	51	25
				0.63	0.100	44	38
				0.59	0.121	37	13
0.60	-123.3	0.67	-2.08	0.66	0.027	31	8
				0.70	0.057	50	17
				0.68	0.069	37	16
1.20	-123.1	0.51	-1.45	0.69	0.028	38	11
				0.74	0.052	52	19
				0.68	0.028	39	18

The first simulation does not use finite element interpolation and serves as a reference. The energy drift was determined over 200,000 steps of each NVE MD simulation by a linear regression. For the static and dynamic properties, values for the inner dihedral are given in the middle of the three rows, outer dihedrals above and below. SD, standard deviation, was determined by the jackknife method (see text). RMS, root mean square.




FIGURE 1. Correlation times of the isomerization dynamics of hexa-argane at different maximum lengths of an element edge used in an MD simulation in the NVE ensemble. (◆): Inner dihedral; (▲, ▼): outer dihedrals. Error bars given by the standard deviation of Table I.

deviation in energy is indicated in Table II; it becomes noticeable only for the biggest elements. The gradient inaccuracy introduced by the interpolation reaches a nonnegligible limit with a maximum edge length of an element of 0.6 rad and larger, although the dynamics did not seem to be affected.

As expected, the performance of the interpolation scheme became better with increasing element

size. However, there seemed to be a best value around 0.6 rad maximum edge length. The slight decline with the largest elements was probably due to the inaccuracy checks. Most of the very large elements failed the test and were not used.

The introduction of the accuracy tests greatly enhanced the accuracy of the interpolation, compared to that found in an earlier study,¹³ where the maximum interpolations per step performance

TABLE II. Finite Elements Interpolation Performance for NVE MD Simulations of 200,000 Steps.

Max. Simplex Edge Length (rad)	Final Numbers of			Interpolations / MD Step Ratio (%)	Interpolations / Explicit Calculation	Largest Energy Difference (kJ / mol)	Inaccurate Gradients (%)
	Interpolations	Vertices	Elements				
0.09	27,890	172,110	405,559	14	0.16	0.00010	0.00
0.15	67,251	132,749	713,235	34	0.51	0.00056	0.00
0.30	92,369	107,631	881,776	46	0.86	0.00185	0.05
0.60	104,196	95,804	589,100	52	1.09	0.01199	4.26
1.20	93,562	106,438	340,126	47	0.88	0.06737	9.72

The interpolations per MD step ratio denotes the performance of the corresponding type of simulation (MC or MD) alone, while the interpolations per explicit calculation ratio gives information about the overall (cumulative) performance, which is important when comparing with simulations with preceding MC grid generation.

of 56% was accompanied by a ratio of inaccurate interpolations of 61% when using 0.6 rad elements.

MD WITH PRECEDING MONTE CARLO (MC)

It was found that an MC simulation,¹⁴ performed prior to the MD simulation was particularly useful in triangulating the conformation space in regions of low energy efficiently and thus led to enhanced performance. The maximum MC step size was chosen to be the maximum length of an element edge to ensure that two consecutive points could always be connected to form an edge. The points that were rejected by the MC criterion were still used as candidate vertices. The irregular pathway of an MC simulation allowed a more efficient tiling than the often straight MD trajectories. In the following simulations 50,000 MC steps were followed by 200,000 steps of MD. For these simulations, the system was weakly coupled¹¹ to an external heat bath at 300 K (NVT simulation). The drift in energy became smaller by doing so (Table III), but the isomerization dynamics were not influenced.

The benefits of the preceding MC simulation are clearly visible when comparing Tables II and IV. Although more points in the conformation space are considered because of the additional MC simulation, the total number of explicit calculations is generally slightly smaller. A value of 0.4 rad is the most favorable maximum element size. The performance is at the maximum (or at least nearly), and the inaccuracy is at a negligible level. Using larger elements, the inaccuracy grows to a critical level, and the performance does not improve any more.

Comparing the statistics of the FE system between simulations with and without preceding MC grid preparations, one can see that the MC grid generations indeed leads to an improved mesh: fewer elements lead to more interpolations. This means that the individual elements are bigger but still allow the same accuracy, because the deviations in energy and forces remain approximately the same.

The preceding MC simulation was carried out at 100 and 300 K for each series of simulations with different maximum lengths of an element

TABLE III.
Energetics and Static and Dynamic Properties of Hexa-Argane Test System.

Max Simplex Edge Length (rad)	Total Energy (kJ/mol)			cis Fraction	Correlation Time (ps)
	Average	Fluctuation	Drift		
0.0	− 133.3	7.3	− 1.95	0.68	42
				0.66	56
				0.66	44
0.2	− 132.9	7.27	+0.476	0.65	41
				0.67	65
				0.64	43
0.3	− 133.1	7.12	+0.004	0.64	43
				0.71	40
				0.65	42
0.4	− 133.2	7.57	+0.477	0.71	40
				0.73	65
				0.69	40
0.5	− 133.2	7.68	+1.144	0.67	33
				0.71	48
				0.64	39
0.6	− 133.1	7.32	+1.079	0.69	34
				0.73	49
				0.67	36

Each NVT MD simulation was 200,000 steps, preceded by 50,000 steps of a 300 K MC grid preparation. Different maximum element edge sizes are shown. The first simulation does not use the finite element interpolation. The drift in energy was determined by a linear regression. For the static and dynamic properties, values for the inner dihedral are given in the middle of the three rows, outer dihedrals above and below.

TABLE IV.
Finite Elements Interpolation Performance for NVT MD Simulations at 300 K with Preceding MC Simulation at 100 or 300 K.

Max. Simplex Edge Length (rad)	Final No.			Interpolations / MD Step Ratio (%)	Interpolations / Explicit Calculations	Largest Energy Difference (kJ / mol)	Inaccurate Gradients (%)
	Interpolations	Vertices	Elements				
100 K MC Temperature							
0.2	88908	120974	915597	44	0.73	0.00098	0.00
0.3	119012	96510	725882	60	1.23	0.00175	0.05
0.4	123432	93132	667353	62	1.33	0.00333	0.31
0.5	125501	91346	586086	63	1.37	0.00614	2.04
0.6	127828	93897	549440	64	1.36	0.01143	6.68
300 K MC Temperature							
0.2	116541	124706	928578	58	0.93	0.00096	0.00
0.3	135947	101786	743803	68	1.34	0.00174	0.06
0.4	143785	95339	657136	72	1.51	0.00329	0.37
0.5	140324	98383	621964	70	1.43	0.00831	1.98
0.6	139777	99482	548806	70	1.41	0.01170	5.68

See also legend of Table II.

edge. The lower MC temperatures allowed a more concentrated sampling of the low energy regions. However, care must be taken that all low energy regions were visited. The 300 K MC simulation did not bias the ensemble of the conformation points

reached but sometimes visited high-energy regions that were not useful in the MD simulation afterward. This observation is confirmed by the results in Table V. A very good performance is obtained in the 100 K MC part. However, the FE system is

TABLE V.
Snapshots of the Time Evolution of Finite Element Interpolation Grid for MD Simulation with Preceding MC Simulation.

No. Steps and Type	No. of			Interpolations / MD	Interpolations / Explicit Calculation
	Interpolations	Vertices	Elements	Step Ratio (%)	
100 K MC Temperature, 0.4 rad Max. Edge Length					
50,00 MC	33,437	16,564	125,993	67	2.02
50,000 MD	27,469	39,095	234,224	55	0.70
100,000 MD	55,034	61,530	371,416	55	0.85
150,000 MD	89,351	77,213	512,881	60	1.16
200,000 MD	123,432	93,132	667,353	62	1.33
300 K MC Temperature, 0.4 rad Max Edge Length					
50,000 MC	10,877	39,124	145,044	22	0.28
50,000 MD	31,581	57,543	287,055	63	0.55
100,000 MD	67,576	71,548	411,961	68	0.94
150,000 MD	103,697	85,427	541,514	69	1.21
200,000 MD	143,785	95,339	657,136	72	1.51

All indicated quantities are cumulative, except the numbers of interpolations which is reset at the start of the MD part. See also legend of Table II.

less well prepared for the follow-up MD part, as indicated by the interpolation per step ratio of the first 50,000 MD steps that is lower than in the example with a 300 K MC preparation. Looking at the overall performance, the 100 K preparation exceeds the results of the 300 K simulation for up to 100,000 steps. This is because the 100 K MC part requires only about one-third of the calculations that the 300 K MC part requires. Different snapshots from the simulations are shown in Table V. The differences are less pronounced for longer simulations. So using a lowered MC temperature is especially effective for shorter MD simulations.

The time evolution of the FE interpolation grid is shown in Figure 2. In the step range less than zero, an MC grid generation took place. The number of interpolations is reset at step zero when the MD started to separate the performance of the two parts. The number of elements grows very quickly and shows only a slight tendency to grow slower. On the other hand, the number of vertices is flattening, which is in agreement with the increasing number of interpolations. The number of interpolations grows nearly linearly, with a slightly increasing trend. Clearly visible is the increased growth in the MD part (larger slope). However, in

spite of the long simulation, saturation (no more new vertices and elements) is not yet reached.

FE MESH HANDLING TIMING PROFILE

To get an idea of how much computer time the different parts of the interpolation method consume, the MC simulation at 300 K with a maximum element edge length of 0.4 rad has been profiled (Table VI). On average, the FE part consumes 13% of the time (including all interpolations, but not explicit calculations), which is affordable. The timings are split up into simulation blocks of 10,000 steps and into the key tasks of the FE mesh handling. In general, the processing time increases as the mesh grows. However, this is not true for the triangulation checks, for which large fluctuations are observed. This is probably due to the current environment of the new points, because it requires more time to establish a correct triangulation in a region already densely populated than in a region visited for the first time. The other quantities do not depend on the surrounding mesh. The increasing time is due to the growing numbers of elements to be considered.

The triangulation checks are the most time consuming, which is in accordance with the complex




FIGURE 2. Time evolution of the finite element interpolation grid. A 300 K MC grid preparation is shown in the negative step range; MD starts at 0. Maximum length of an element edge is 0.4 rad.

TABLE VI.
Timings of Different Parts of the FE Mesh Handling on Sun Ultra 1 (200 MHz).

Description	Steps of Simulation Block					
	1–10,000	10,001–20,000	20,001–30,000	30,001–40,000	40,001–50,000	1–50,000
Absolute Timings (s)						
Handling total	237.43	903.21	1733.22	1225.62	2553.70	6653.18
Find element	67.02	191.26	273.44	333.91	400.49	1266.12
Build elements total	153.76	651.01	1359.09	754.58	2153.58	5072.02
Triangulation checks	67.96	366.34	888.32	93.09	1520.70	2936.41
Element properties checks	50.59	173.16	287.42	395.67	404.12	1310.96
Element accuracy checks	30.74	95.26	156.84	219.94	203.92	706.70
Time / Vertex in Mesh (s)						
Handling total	0.0248	0.0497	0.0674	0.0377	0.0653	0.1701
Find element	0.0070	0.0105	0.0106	0.0103	0.0102	0.0324
Build elements total	0.0161	0.0358	0.0529	0.0232	0.0550	0.1296
Triangulation checks	0.0071	0.0201	0.0345	0.0029	0.0389	0.0751
Element properties checks	0.0053	0.0095	0.0112	0.0122	0.0103	0.0335
Element accuracy checks	0.0032	0.0052	0.0061	0.0068	0.0052	0.0181
Vertices	9,575	18,183	25,714	32,519	39,124	39,124
Elements	11,555	37,639	69,639	107,962	145,044	145,044

An MC simulation of 50,000 steps at 300 K and a maximum element edge length of 0.4 rad is analyzed in five blocks of 10,000 steps and in total, shown in the six columns. Handling total: time required for the finite element mesh handling in total, which consists roughly of the two following rows. Find element: time required to scan the existing elements for an element that contains the actual point (for interpolation). Build elements total: time required in total to build new elements if no interpolation took place. This quantity is again roughly split up in the next three filters. Triangulation checks: time required to ensure correct triangulated mesh. Element properties checks: time required to check edge lengths and angles of a candidate element. Element accuracy checks: time required to verify the interpolation accuracy in the candidate element (includes interpolation at every vertex of the element). Vertices Elements: number of vertices and elements in the mesh after the corresponding simulation block.

calculations that are involved. The accuracy checks are relatively cheap but enable a much more efficient interpolation (see above). The time required to find the element enclosing a point is nicely proportional to the number of vertices, which is clearly a consequence of the way of finding it by scanning the vertices.

Conclusions

An interpolation scheme has been developed with the enhancement of efficiency of combined quantum-chemical/classical MD simulations (hybrid models) in mind. The method is based on FE interpolation, which is able to treat the irregular nature of the triangulation of the conformation space of a molecule that takes part in a chemical reaction. Irregular tetrahedra were used as FEs. The method was tested for a 3-dimensional case with a simple system. Although the program was implemented for multiple dimensions, 4- and

higher dimensional cases were neither tested nor studied. However, no fundamental problems should arise when generalizing to more dimensions. The method was shown to work properly and efficiently. No significant deviations in static and dynamic properties were observed between simulations with and without interpolations. However, in explicit comparisons of forces, inaccurate interpolations were no longer negligible if the maximum edge lengths of simplex elements became larger than 0.4 rad. Using the largest maximum edge lengths that guarantee accurate forces, nearly 50% of the required (quantum) calculations could be replaced by cheap interpolations in pure MD simulations. This corresponded to an interpolation per explicit calculation ratio of about 0.9. The performance could be further enhanced by preceding the MD simulation by an MC mesh generation. Then the interpolation per explicit calculation ratio could be increased to 1.5 (more than 70% interpolations in the MD part). For short MD simulations, a reduc-

tion of the MC acceptance temperature to one-third of the MD temperature seemed useful.

The presented interpolation method is a promising tool to lower the computational costs of combined QM/classical dynamics simulations by the reduction of explicit QM energy and force evaluations. Work is in progress for a further refinement of the presented method to allow even better performance (e.g., by using all information at the vertices instead of a reduced set) and to incorporate the method into a hybrid calculation.

References

1. A. Warshel and M. Levitt, *J. Mol. Biol.*, **103**, 227 (1976).
2. M. J. Field, P. A. Bash, and M. Karplus, *J. Comput. Chem.*, **11**, 700 (1990).
3. H. Liu, F. Müller-Plathe, and W. F. van Gunsteren, *J. Chem. Phys.*, **102**, 1722 (1995).
4. H. Liu, F. Müller-Plathe, and W. F. van Gunsteren, *Chemistry*, **2**, 191 (1996).
5. W. F. van Gunsteren and H. J. C. Berendsen, *Angew. Chem. Int. Ed. Engl.*, **29**, 992 (1990).
6. M. P. Allen and D. J. Tildesley, *Computer Simulation of Liquids*, Oxford Science Publications, Oxford, U.K., 1987.
7. F. Müller-Plathe, unpublished results.
8. P. G. Ciarlet, In *Handbook of Numerical Analysis*, Vol. II, Part 1, P. G. Ciarlet and J. L. Lions, Eds., Elsevier, Amsterdam, 1991.
9. E. B. Becker, G. F. Carey, and J. T. Oden, *Finite Elements*, Prentice-Hall, Englewood Cliffs, NJ, 1986.
10. J.-P. Ryckaert, G. Ciccotti, and H. J. C. Berendsen, *J. Comput. Phys.*, **23**, 327 (1977).
11. H. J. C. Berendsen, J. P. M. Postma, W. F. van Gunsteren, A. DiNola, and J. R. Haak, *J. Chem. Phys.*, **81**, 3617 (1984).
12. B. Efron, *The Jackknife, the Bootstrap, and Other Resampling Plans*, SIAM, Philadelphia, PA, 1982.
13. C. D. Berweger, Diplour thesis, ETH, Zürich, 1995.
14. N. Metropolis and S. Ulam, *J. Am. Stat. Assoc.*, **44**, 335 (1949).